



AFRL-OSR-VA-TR-2013-0578

**DATA-DRIVEN ONLINE AND REAL-TIME COMBINATORIAL
OPTIMIZATION**

PATRICK JAILLET

MASSACHUSETTS INSTITUTE OF

10/30/2013

Final Report

DISTRIBUTION A: Distribution approved for public release.

**AIR FORCE RESEARCH LABORATORY
AF OFFICE OF SCIENTIFIC RESEARCH (AFOSR)/RSL
ARLINGTON, VIRGINIA 22203
AIR FORCE MATERIEL COMMAND**

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188		
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 28-10-2013		2. REPORT TYPE Final		3. DATES COVERED (From - To) From 01-09-2010 To 31-07-2013	
4. TITLE AND SUBTITLE Data-Driven Online and Real-Time Combinatorial Optimization			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER FA9550-10-0437		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Jaillet, Patrick			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology Office of Sponsored Programs 77 Massachusetts Ave Cambridge, MA 02139-4301			8. PERFORMING ORGANIZATION REPORT NUMBER 6922649		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research 875 N. Randolph Street Arlington, VA 22203-1954			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The main focus of our research has been on the fundamental aspects of optimization in the context of uncertain (and possibly large) data sets revealed in an online fashion, considering the intersection and interplay of three main phenomena (incomplete and uncertain data, online decisions with or without real-time restrictions, and large data sets). Motivated by applications associated with the deployment of autonomous multi-agent systems for spatial exploration and information harvesting, our research has concentrated on the development and analysis of competitive online algorithms for the simplest canonical models defined in our proposal (single agent prize collecting online traveling salesman problem and Hamiltonian path problems), as well as for some generalizations of the secretary problem, a class of closely related online problems.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

DATA-DRIVEN ONLINE AND REAL-TIME COMBINATORIAL OPTIMIZATION

Patrick Jaillet, MIT

AFOSR grant FA9550-10-1-0437

Program Director - Dr. Donald Hearn

Final Report, October 2013

1 Overview of the PI's long-range research objectives

The main focus of our research is on the fundamental aspects of optimization in the context of uncertain (and possibly large) data sets revealed in an online fashion. More specifically, we are interested in the intersection and interplay of three main phenomena (incomplete and uncertain data, online decisions with or without real-time restrictions, large data sets), and the corresponding fundamental questions when facing a problem exhibiting one or several of these phenomena.

There are many motivating examples for such problems. In the proposal that led to this funded research, we highlighted applications associated with the deployment of autonomous multi-agent systems for spatial exploration and information harvesting.

The major objectives of this long-ranging research involve the definition and rigorous mathematical analysis of canonical models which would capture the essence of this new class of problems. The overall purpose is to systematically (i) analyze and integrate the most promising paradigms/solution strategies for dealing with these problems, (ii) formulate adequate corresponding models, and (iii) develop, analyze, and implement algorithms for solving them.

2 Results from grant FA9550-10-1-0437

2.1 Research emphasis

Our research has concentrated on the development and analysis of online algorithms on some of the simplest canonical models defined in our proposal (restricting our research to single agent prize collecting online TSP and Hamiltonian path problems). We have also considered generalizations of the secretary problem, a class of closely related online problems.

2.2 Technical details of results

1. Lehouillier, T. “Online TSP with Moving Targets”, Research Project Final Report, supervisor, P. Jaillet, EECS MIT and ENSIMAG Grenoble, September 2011.

Among other results this report concentrates on the online TSP with moving targets: Given S a set of n requests s_i , having a release date r_i , an initial position l_i , and a constant speed \vec{v}_i , and given pursuer starting at the origin having maximum speed v such that $v > |\vec{v}_i|$, find the shortest tour ending at the origin after serving all requests. We first restrict ourselves to problems on the non-negative line for which all requests have the same speed

v_R and the server a speed $v_S > v_R$. We show that every c -competitive online algorithm has $c \geq \frac{3}{2} + \frac{v_R}{v_S - v_R}$. We then propose an algorithm ("Move Right if Necessary") achieving this competitive ratio. We then consider the problem on the real line and show that any c -competitive on-line algorithm for the on-line TSP with moving requests on the real line has $c \geq \frac{7}{4} + \alpha(3 + 2\alpha) - \frac{\alpha}{4(1 + \alpha)}$, where $\alpha = \frac{v_R}{v_S - v_R}$.

2. Zenklusen, R., "Randomized Competitive Ratio of TSP on Non-negative Line", technical report, MIT, March 2012.

Many hardness results for competitive ratios are based on the assumption that the online algorithm is deterministic. What happens if this algorithm is allowed to be randomized? It seems interesting to study for a given online problem, how much one can gain through randomization. In general, randomization might be one natural way to fight against an overly powerful adversary. Here we prove that even when allowing randomization in the algorithm, it is not possible to obtain an algorithm for the online TSP on the non-negative real line with competitive ratio $< 1.5 - \varepsilon$ for any $\varepsilon > 0$.

3. Hwang D., "Online Routing Problems", Report, EECS, MIT, Nov. 2012.

This report focuses on online routing problems, including the online Hamiltonian Path Problem, the online Traveling Salesman Problem, and variations of the online Quota Hamiltonian Path Problem and the online Traveling Repairman Problem.

First, we consider the online Hamiltonian Path Problem (HPP). In this problem, the objective is to minimize the makespan, i.e., the earliest time when all requests are visited. When the metric space is the non-negative real line, we provide a novel online algorithm whose competitive ratio is 1.75, lower than the best existing one (2.06). The algorithm only keeps track of the left-most and right-most outstanding requests, and it only re-calculates the route at the arrival of a new request that becomes the new left-most outstanding request. We also show that the proposed algorithm has the lowest competitive ratio among all algorithms of this kind.

Second, we consider the Online Traveling Salesman Problem, and consider randomized algorithms. We show that the competitive ratio of a given randomized online algorithm is smaller than 1.93 when the metric space is the discrete star graph of any size. Combined with the result in the literature that the competitive ratio of any deterministic algorithm approaches 2 as the size of the star graph goes to infinity, we conclude that randomized algorithms can reach lower competitive ratios than its deterministic counterparts for the online TSP on some metric spaces.

Third, we consider the Online Quota Hamiltonian Path Problem (QHPP) with knowledge of release dates. In the online QHPP, each request has a weight that is revealed with the request. The goal is to minimize the time at which the sum of the weight of visited requests reaches a certain prescribed quota Q . We consider the case in which online algorithms have the same knowledge of the release dates as offline algorithms do, but online algorithms only know the location and weight of a request at its release date while offline algorithms know this information at time 0. We consider two models of the release dates. In the first model, both online and offline algorithms know the precise release dates. In the second model, both online and offline algorithms know only the distribution of the release dates. We obtain the

best-possible competitive ratios in some simple cases.

Fourth, we consider the online QHPP with noisy weight. In this problem, both online and offline algorithms observe a noisy weight, assumed to be the true weight plus an i.i.d. Gaussian noise. Neither the online nor the offline algorithms can observe the true weight of a request even after visiting it. The objective is to minimize the time at which the probability that the sum of the true weight of visited requests reaches the prescribed quota Q is at least a given threshold p . The on-site service time can be learned, but is only known by any algorithm when it is completed. We provide an algorithm with a competitive ratio of 2, and we prove that 2 is best-possible.

Fifth, we consider the online QHPP with uncertain service time. In this problem, we consider the case in which the service of a request is completed when the algorithm visits the request for a consecutive duration equal to the required on-site service time, specified by the request. Algorithms are allowed to leave the location of a request that is being served. However, if it decides to come back and serve the request later, it has to start-over from scratch. We consider the case in which both online and offline algorithms do not know the precise required on-site service time, but only its distribution, at the time the request is revealed. We propose online algorithms with finite competitive ratios for this problem.

Sixth, we consider the online Weighted Traveling Repairman Problem (WTRP) with uncertain service time. The notion of the uncertainty of service time is the same as described in the fifth problem. In this problem, the objective is to minimize the sum of the weighted completion time, where the completion time of a request is defined to be the time when the service of the request is completed. We propose online algorithms with finite competitive ratios for this problem.

4. Gupta, S., “Discrete Online TSP”, Report, ORC, MIT, April 2013.

We give a 3-competitive polynomial time algorithm for nomadic online TSP on undirected edge-weighted graphs. Requests arrive on the vertices of the given graph and can be served only after their release time. We further restrict the server to not re-route on an edge. This algorithm generalizes to metric spaces.

5. Jaillet, P., J. Soto and R. Zenklusen, “Advances on Matroid Secretary Problems: Free Order Model and Laminar Case”, accepted for publication in the proceedings of the 16th Conference on Integer Programming and Combinatorial Optimization (IPCO 2013), January 2013 (was initially submitted July 2012). An extended journal version is in final preparation for submission to Mathematical Programming.

The best-known conjecture in the context of matroid secretary problems claims the existence of an $O(1)$ -approximation applicable to any matroid. Whereas this conjecture remains open, modified forms of it have been shown to be true, when assuming that the assignment of weights to the secretaries is not adversarial but uniformly at random. However, so far, no variant of the matroid secretary problem with adversarial weight assignment is known that admits an $O(1)$ -approximation. We address this point by presenting a 9-approximation for the free order model, a model suggested shortly after the introduction of the matroid secretary problem, and for which no $O(1)$ -approximation was known so far. The free order model is a relaxed version of the original matroid secretary problem, with the only difference being that one can choose the order in which secretaries are interviewed.

Furthermore, we consider the classical matroid secretary problem for the special case of

laminar matroids. Only recently, a $O(1)$ -approximation has been found for this case, using a clever but rather involved method and analysis that leads to a $16000/3$ -approximation. This is arguably one of the most involved special case of the matroid secretary problem for which an $O(1)$ -approximation is known. We present a considerably simpler and stronger $33e^{14.12}$ -approximation, based on reducing the problem to a matroid secretary problem on a partition matroid.

6. Jaillet, P. and X. Lu. “Online Traveling Salesman Problems with Rejection Options”, submitted to *Networks*, March 2013; invited to revise and resubmit, September 2013.

In this paper we consider online versions of the Traveling Salesman Problem (TSP) on metric spaces for which requests to visit points are not mandatory. Associated with each request is a penalty (if rejected). Requests are revealed over time (at their release dates) to a server who must decide which requests to accept and serve in order to minimize a linear combination of the time to serve all accepted requests and the total penalties of all rejected requests. In the *basic* online version of the problem, a request can be accepted any time after its release date. In the *real-time* online version, a request must be accepted or rejected at the time of its release date.

For the basic version, we provide a best possible 2-competitive online algorithm for the problem on a general metric space. For the real-time version, we first consider special metric spaces: on the non-negative real line, we provide a best possible 2.5-competitive polynomial time online algorithm; on the real line, we prove a lower bound of 2.64 on any competitive ratios and give a 3-competitive online algorithm. We then consider the case of a general metric space and prove a $\Omega(\sqrt{\ln n})$ lower bound on the competitive ratio of any online algorithms. Finally, among the restricted class of online algorithms with prior knowledge about the total number of requests n , we propose an asymptotically best possible $O(\sqrt{\ln n})$ -competitive algorithm.

7. Hwang D. and P. Jaillet, “Better Algorithms for Online Traveling Repairman Problems”, Working paper, EECS, MIT, July 2013.

In this paper, we consider the online Weighted Traveling Repairman Problem (WTRP) in which a single repairman with a unit speed limit is to travel with the objective of minimizing the total weighted completion time of a finite number of requests in any general metric space. We consider the online time model in which the requests are revealed over time, and at any given time, online algorithms only know the partially revealed problem instance. No probabilistic assumptions on the release dates, the locations, or even the number of requests are imposed. We propose and analyze both deterministic and randomized online algorithms, proving that our algorithms are better than the existing ones with respect to competitive ratios, the standard performance measure of online algorithms. In addition, we also test the performance of our algorithm in randomly generated problem instances. For deterministic online algorithms, we propose an algorithm whose competitive ratio is between 4 and 5.14 in general metric spaces. The upper bound is lower than the best existing one, which is $3 + 2\sqrt{2} \approx 5.83$. For randomized online algorithms, we propose an algorithm whose competitive ratio is between $1 + 2/\ln(3) \approx 2.82$ and $4/\ln(3) \approx 3.64$ in general metric spaces. The upper bound is lower than the best existing one, which is $(2 + \sqrt{2})/\ln(1 + \sqrt{2}) \approx 3.87$. The upper bound on the competitive ratios are also valid for many other online vehicle routing problems when the objective is to minimize the total weighted completion time.

Those problems include the online dial-a-ride problem and the online TRP with precedence constraints.

2.3 Funded personnel and facilities

Personnel: The personnel funded by this grant has included:

Prof. Jaillet, PI: partial summer salary support over the grant duration.

Dawsen Hwang, MIT EECS doctoral student: fully supported (RA stipend + tuition) since January 15, 2011. Tasks:

Swati Gupta, MIT ORC doctoral student: partially supported (1/2 RA stipend + tuition) since 01-09-2011. Tasks:

Thibault Lehouillier, visiting student from France (ENSIMAG Grenoble), completed a 6 months research internship (March to August 2011).

Dr. Rico Zenklusen, MIT post-doc, partially funded on this project from August 1, 2011 until July 31, 2012.

Other personnel working on this research without direct funding included:

Prof. Goemans, Math, MIT. Participated in research group meetings with Swati Gupta, as her co-supervisor

Xin Lu, MIT ORC doctoral student. Participated in various research group meetings.

Facilities: The main facilities available for conducting this research included offices and computing capabilities within the Laboratory for Information and Decision Systems (LIDS) and the Operations Research Center (ORC) at MIT.

2.4 Presentations and other information

Presentations: Our research has been presented at major national and international conferences (such as INFORMS annual meetings, TRISTAN international meetings, Optimization Days) as well as in select specialized workshops (AFOSR-DDDAS, Washington; ONR-IPAM, UCLA). The research has also been presented at invited seminars at various universities (CISE, Boston University, LIDS, MIT).

Honors and recognitions: Over the duration of this grant, the PI received the Dugald C. Jackson Professorship in the Department of Electrical Engineering and Computer Science, MIT; became Co-Director of the MIT Operations Research Center; received in 2012 the 2010 Glover-Klingman Prize; and was recently inducted as an INFORMS Fellow with the following citation: “for opening the field of a priori optimization for stochastic programming and for being regarded as the world leader in the field of probabilistic and online optimization”

Derived outcomes: From conference presentations and workshop presentations, as described above, contacts have been made which have resulted in the following derived activities:

Co-supervision of an MIT SM ORC thesis of a Draper Fellow from the Navy, Jacob Cates, entitled “Route Optimization Under Uncertainty for Unmanned Underwater Vehicles”.

Collaboration with a team from Lockheed Martin Advanced Technology Laboratories in the writing of a white paper and forthcoming participation as a consultant in a recently funded proposal from Lockheed Martin ATL in response to ONR SN 11-SN-0012 (PI Ray Yuan: Probabilistic Auction for Distributed, Robust, Resource, Optimization).

3 Conclusion

In our research we have successfully looked at some of the fundamental intellectual challenges behind (1) the integration of (i) incomplete and uncertain data, (ii) short time restriction for decisions, (iii) large data sets; (2) the development of appropriate canonical models for considering these multi-facet issues, and (3) the design of new algorithmic solutions and paradigms for addressing these canonical models. Due to the richness of the proposed set of problems, we have restricted our analysis to single-agent systems. Future opportunities and challenges occur when considering multi-agent systems operating under these conditions in a communication limited environment. We are currently pursuing this line of research and have recently obtained results on distributed online assignment for multi-depot vehicle routing problems.